

What is claimed is:

1. A method of controlling power consumption for at least one computer system,  
wherein a power supply for the at least one computer system has a maximum power  
5 output based on a nominal power consumption of the at least one computer system,  
the method comprising:  
detecting an amount of power consumed by the at least one computer system;  
comparing the amount of power consumed by the at least one computer system  
to a threshold, wherein the threshold is based on the maximum power output of the  
10 power supply; and  
placing one or more components of the at least one computer system in a  
lower-power state to reduce power consumption in response to the amount of power  
consumed by the at least one computer system exceeding the threshold.
- 15 2. The method of claim 1, wherein a cooling system is operable to cool the at  
least one computer system, the method further comprising:  
determining whether insufficient cooling resources are available for cooling  
the at least one computer system; and  
placing at least one component of the at least one computer system in a lower-  
20 power state in response to insufficient cooling resources being available to cool the at  
least one computer system.

3. The method of claim 2, further comprising:

determining whether excess cooling resources are available for cooling the at least one computer system; and

placing the at least one component of the computer system currently in a lower-power state in a higher-power state, such that the at least one component consumes more power, in response to excess cooling resources being available.

4. The method of claim 3, wherein determining whether excess cooling resources

are available for cooling the at least one computer system comprises determining

whether an amount of cooling fluid distributed to the at least one computer system is less than an excess cooling fluid threshold.

5. The method of claim 2, wherein determining whether insufficient cooling resources are available for cooling the at least one computer system comprises

determining whether an amount of cooling fluid distributed to the at least one computer system exceeds a threshold associated with the maximum capacity of the cooling system.

6. The method of claim 1, further comprising:

comparing the amount of power consumed by the at least one computer system to a second threshold;

placing the at least one component of the computer system, currently in a lower-power state, in a higher-power state, such that the at least one component

consumes more power, in response to the amount of power consumed by the at least one computer system being less than the second threshold.

7. The method of claim 6, wherein placing the at least one component of the computer system in a higher-power state comprises:

determining whether placing the at least one component in a higher-power state will cause the power consumption of the at least one computer system to exceed the threshold based on the maximum power consumption of the at least one computer system; and

placing the at least component in a higher-power state in response to determining the power consumption will not exceed the threshold.

8. The method of claim 6, wherein one or more of the threshold based on the maximum power output of the power supply and the second threshold is determined such that a minimal change in power consumption does not result in changing a power state of the at least one component.

9. The method of claim 1, wherein placing the at least one component in a lower-power state comprises:

determining a cooling efficiency of components in the at least one computer system; and

selecting one or more of the components to be placed in a lower-power state based on an amount of energy needed to cool the one or more components; wherein a

component requiring more energy to be cooled is selected before a component requiring less energy to be cooled.

10. The method of claim 1, wherein the at least one computer system comprises  
5 multiple computer systems running applications, the method further comprising:  
prioritizing applications running on the multiple computer systems; wherein  
the step of placing one or more components in a lower-power state further  
comprises identifying one of the multiple computer systems running one or more low  
priority applications, and placing at least one component in the identified computer  
10 system in a lower-power state.

11. The method of claim 1, wherein a processor for the at least one computer  
system is operable to be placed in multiple lower-power states, each lower-power state  
being associated with a lower clock speed, and placing one or more components of the  
15 at least one computer system in a lower-power state comprises placing the processor  
in one of the multiple lower-power states.

12. The method of claim 11, wherein placing the processor in one of the multiple  
lower-power states comprises instructing the processor not to consume more than a  
20 predetermined amount of power.

13. The method of claim 1, further comprising:

storing information including components in the at least one computer system,  
power state of the components, power consumption of the components, and priority  
information associated with prioritizing components placed in a lower-power state;

and placing one or more components of the at least one computer system in a  
5 lower-power state comprises determining the one or more components to be placed in  
a lower-power state based on the stored information.

14. The method of claim 13, further comprising:

placing the at least one component of the computer system currently in a  
10 lower-power state in a higher-power state based on the stored information.

15. The method of claim 1, wherein placing one or more components in a low-  
power state comprises reducing power consumption of one or more of a processor, a  
floating point unit, one or more storage devices, main memory, and a cache or a  
15 portion of a cache.

16. A power system generating power for at least one computer system, the power  
system comprising:

at least one power supply operable to provide power for the at least one  
20 computer system;

a power monitor operable to determine the power consumption of the at least  
one computer system; and

a power provisioning system operable to compare the power consumption of  
the at least one computer system to a threshold associated with a maximum capacity

of the power supply, and further operable to place one or more components of the at least one computer system in a lower-power state in response to the power consumption exceeding the threshold.

5        17.     The power system of claim 16, wherein the power supply is designed based on a nominal power consumption of the at least one computer system.

18.     The power system of claim 17, wherein the maximum capacity of the at least one power supply is approximately equal to the nominal power consumption of the at least one computer system.

10

19.     The power system of claim 17, wherein the power provisioning system is connected to a cooling system and is operable to receive messages from the cooling system associated with the availability of cooling resources for cooling the at least one computer system, the power provisioning system being operable to control the power consumption of the at least one computer system based on a message received from the cooling system.

15

20.     The power system of claim 19, wherein the power provisioning system is operable to place at least one component of the at least one computer system in a lower-power state in response to receiving a message from the cooling system indicating that insufficient cooling resources are available for cooling the at least one computer system.

20

21. The power system of claim 20, wherein the power provisioning system is operable to place at least one of the at least one component currently in a lower-power state in a higher-power state in response to receiving a message from the cooling system indicating that excess cooling resources are available.

5

22. The power system of claim 16, wherein the one or more components comprise a processor, and the power provisioning system is operable to instruct the processor to reduce clock speed for reducing power consumption.

10

23. The power system of claim 16, wherein the one or more components comprises a processor, and the power provisioning system is operable to instruct the processor to reduce power consumption of the processor to a calculated value or range of values.

15

24. The power system of claim 16, wherein the one or more components comprises a processor operable to be placed in one of multiple lower-power states.

20

25. The power system of claim 16, wherein the at least one computer system comprises multiple computer systems, and the power provisioning system is operable to prioritize the multiple computer systems for placement in a lower-power state based on an importance of applications executing on the multiple computer systems.

26. The power system of claim 16, further comprising a repository storing power state information for the one or more of components in the at least one computer

system, wherein the power provisioning system is operable to utilize the power state information to identify a component of the one or more components to be placed in a lower-power state or a higher-power state.

5           27.     The power system of claim 26, wherein the power state information comprises one or more of power consumption of the one or more components and priority information associated with prioritizing the one or more components for changing the power state of the one or more components.

10          28.     The power system of claim 16, wherein the power provisioning system is operable to determine whether the amount of power consumed by the at least one computer system falls below a second threshold, and further operable to place at least one of the one or more components in a higher-power state in response to the amount of power consumed by the at least one computer system falling below the second  
15          threshold.

          29.     The power system of claim 16, wherein the at least one computer system comprises multiple computer systems receiving power via a power bus, and the power provisioning system is operable to disconnect a portion of a power bus to place one of  
20          the multiple computer systems in a lower-power state.

          30.     The power system of claim 16, wherein the power monitor is connected to the at least one power supply to measure the output power of the at least one power supply for determining the power consumption of the at least one computer system.



31. The power system of claim 16, wherein the at least one computer system comprises multiple computer systems connected to the at least one power supply via a power bus, and the power monitor is connected to the power bus to measure the power consumption of the multiple computer systems.

5

32. The power system of claim 16, wherein the one or more components comprise one or more of a processor, a floating point unit, one or more storage devices, one or more memory ICs, and a cache or a portion of a cache.

10

33. A system comprising:

multiple computers housed in an enclosure;

a cooling system operable to distribute cooling fluid to the multiple computer systems in the enclosure based on one or more of the power consumption and heat dissipation of the multiple computer systems; and

15

a power system connected to the cooling system and including a power supply operable to generate power for the multiple computer systems and a power provisioning system, wherein the power provisioning system is operable to control power consumption of at least one of the multiple computer systems based on an availability of cooling resources for cooling the multiple computer systems.

20

34. The system of claim 33, wherein one or more of (1) the cooling system is designed based on a nominal heat dissipation of the multiple computer systems and (2) the power supply is designed based on the nominal power consumption of the multiple computer systems.

35. The system of claim 34, wherein the cooling system is operable to transmit a message to the power provisioning system indicating insufficient cooling resources are available for cooling the multiple computer systems or excess cooling resources are available for cooling the multiple computer systems; and

the power provisioning system is operable to reduce power consumption of at least one of the multiple computer systems in response to receiving a message indicating insufficient cooling resources are available for cooling the multiple computer systems; and

the power provisioning system is operable to increase power consumption of at least one of the multiple computer systems in response to receiving a message indicating excess cooling resources are available for cooling the multiple computer systems.

36. The system of claim 34, wherein the power provisioning system is operable to compare the power consumption of the multiple computer systems to a threshold associated with a maximum capacity of the power supply and reduce the power consumption of at least one of the multiple computer systems in response to the power consumption exceeding the threshold.

20

37. The system of claim 33, wherein the enclosure is a rack.

38. The system of claim 33, wherein the enclosure is a data center.

39. An apparatus controlling power consumption of at least one computer system using a power supply means having a maximum power output based on a nominal power consumption of the computer system, the apparatus comprising:

means for determining an amount of power consumed by the at least one computer system;

means for comparing the amount of power to a threshold, wherein the threshold is based on the maximum power output of the power supply means; and

means for placing one or more components of the at least one computer system in a lower-power state to reduce power consumption in response to the power consumption of the at least one computer system exceeding the threshold.

40. The apparatus of claim 39, further comprising a means for cooling the at least one computer system, wherein the means for cooling is further operable to determine whether insufficient cooling resources are available for cooling the at least one computer system and the means for placing at least one component of the at least one computer system in a lower-power state is further operable to place the at least one component in a lower-power state in response to insufficient cooling resources being available to cool the at least one computer system.

41. The apparatus of claim 39, wherein the means for cooling is further operable to determine whether excess cooling resources are available for cooling the at least one computer system; and

the means for placing the at least one component of the computer system currently in a lower-power state is further operable to place the at least one component

in a higher-power state, such that the at least one component consumes more power,  
in response to excess cooling resources being available.

42. The apparatus of claim 39, wherein the means for comparing the amount of  
5 power consumed by the at least one computer system to a threshold is further operable  
to compare the amount of power consumed by the at least one computer system to a  
second threshold;

and the means for placing the at least one component of the computer system  
in a lower-power state is further operable to place the at least one component in a  
10 higher-power state, such that the at least one component consumes more power, in  
response to the amount of power consumed by the at least one computer system being  
less than the second threshold.

43. A method of designing at least one computer system comprising:  
15 selecting components for a computer system;  
selecting a power supply for the computer system based on a nominal power  
consumption of the components being used in the computer system, the nominal  
power consumption being a power consumption less than a maximum power  
consumption of the components being used in the computer system; and  
20 assembling the computer system, the computer system including the selected  
components and the selected power supply.

44. The method of claim 43, wherein the computer system comprises a rack  
system housing multiple computer systems, the method further comprising:

determining a number of computer systems operable to be powered by the selected power supply based on a nominal power consumption of each computer system; and

installing the number of computer systems in the rack, such that the number of computer systems are powered by the selected power supply.

45. The method of claim 43, wherein the nominal power consumption is based on an average power consumption of the components when used in the computer system.

46. The method of claim 43, wherein the components comprise one or more of at least one processor, memory, storage device, and cooling system.